

Causal Modeling and Event-driven Simulation for Monitoring of Continuous Systems

Richard Doyle, Leonard Charest, Jr., Nicolas Rouquette and Jay Wyatt

Jet Propulsion Laboratory
Pasadena, CA 91109-8099

Charles Robertson
McDonnell-Douglas Aerospace
Houston, TX 77059-3556

Abstract

Monitoring of complex continuous physical systems has been traditionally accomplished in computer-based process control software by one or both of the following methods: 1) establishing limit checks for sensors and raising an alarm whenever a sensor's value crosses one of these thresholds, and 2) comparing predicted values from a simulation to actual sensor values and flagging discrepancies. These anomaly detection techniques are not as robust as they need to be. Failures can manifest in ways which are not captured by these traditional methods. Furthermore, some anomalous behaviors are more naturally detected at the level of global interactions affecting multiple sensors.

We describe extensions to the traditional techniques for anomaly detection, as well as new anomaly detection techniques based on alternate models of what distinguishes normal from abnormal behavior. Some of these techniques are designed to capture anomalies at individual sensors; some detect anomalies across collections of sensors. To assist in reasoning about complex global behaviors, we construct and simulate a causal model of the physical system being monitored.

These techniques have been tested on data from the Environmental Control and Life Support System (ECLSS) of Space Station Freedom (SSF) and are being applied in advanced monitoring prototypes for the SSF External Active Thermal Control System (EATCS) of SSF and the Environmental Emergency and Consumables Management (EECOM) subsystem of the Space Shuttle.

1 Introduction

Mission operations personnel at NASA have the task of determining, from moment to moment, whether a space platform is exhibiting behavior which is in any way anomalous, which could disrupt the operation of the platform, and in the worst case, could represent a loss of ability to achieve mission goals. A traditional technique for assisting mission operators in space platform health analysis is the establishment of alarm thresholds for sensors, typically indexed by operating mode, which summarize which ranges of sensor values imply the

existence of anomalies. Another established technique for anomaly detection is the comparison of predicted values from a simulation to actual values received in telemetry. However, experienced mission operators reason about more than alarm threshold crossings and discrepancies between predicted and actual to detect anomalies: they may ask whether a sensor is behaving differently than it has in the past, whether a current behavior may lead to a global perturbation or whether a current behavior may lead to the particular bane of operators - a rapidly developing alarm sequence.

A fault which propagates through a system faster than the sensor polling rate can create a situation where, between one sampling and the next, the number of sensors in alarm goes from zero to tens or more. Information about the ordering of events is lost. In this kind of emergency situation, operators can experience information overload and a compromising of their ability to interpret the sensor data.

Our approach to introducing automation into real-time systems monitoring is based on two observations: 1) mission operators employ multiple methods for recognizing anomalies, and 2) mission operators do not and should not interpret all sensor data all of the time. The subject of this paper is an approach to determining from moment to moment which subset of the available sensor data for a system is most informative about the presence of, or potential for, anomalies occurring within the system. We term this process *sensor selection* and we have implemented a prototype selective monitoring system called SELMON [7; 8].

The SELMON system has its origins in a sensor planning system called GRIPE [7] which planned information gathering activities to verify the execution of robot task plans. Other model-based monitoring systems include Dvorak's MIMIC, which performs robust discrepancy detection for continuous dynamic systems [9; 10], and DeCoste's DATMI, which infers system states from incomplete sensor data [5]. The SELMON work complements other work within NASA which has focused on empirical and model-based methods for fault diagnosis of aerospace platforms [1; 13; 11].

The organization of this paper is as follows: First we describe sensor importance measures used to identify the presence of anomalies in a monitored system, and then appraise the relevance of different sensors for reporting on those anomalies. Two of these measures are extensions of the traditional techniques of limit sensing and discrepancy detection. The

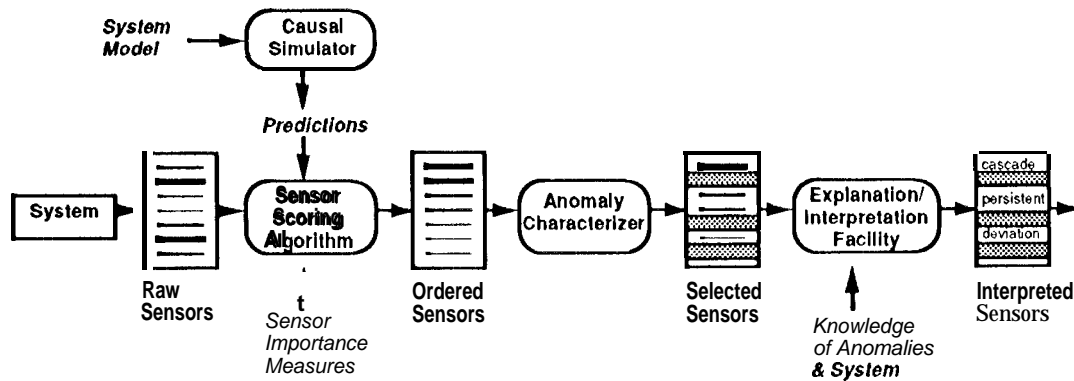


Figure 1: SELMON Architecture.

seven measures fall into three categories: those which concern themselves with data provided by single sensors, those which utilize a simulation of the system being monitored, and those which utilize a causal model of the system being monitored. As part of our discussion of the use of causal modeling and reasoning, we briefly describe our event-driven causal simulator and our causal modeling language.

Next we describe the test and application domains we are working with, providing more detail on the SSF/EATCS. Then we report and discuss results on empirically evaluating the performance of the SELMON system on data from the SSF/EATCS testbed. As part of this section, we describe how SELMON detected an anomaly which the traditional monitoring methods fail to detect. Finally, we conclude with some thoughts on future research and applications and a summary.

2 Approach: Selective Monitoring

How does an intelligent agent – human or machine – observing a complex physical system, decide when something is going wrong? Abnormal behavior is always defined as a departure from normal behavior. Unfortunately, there appears to be no single, crisp definition of “normal” behavior. In the traditional monitoring technique of limit sensing, normal behavior is predefined by nominal value ranges for sensors. A fundamental limitation of this approach is the lack of sensitivity to context. A single fault may manifest in different ways, depending on the configuration of the system when the fault occurs. The compiled notion of an alarm threshold may not capture these subtleties in manifestation. Another limitation is the purely local view of sensors taken by the limit sensing technique. It is possible for a system to exhibit a pattern of sensor values, each value in its respective nominal range, which is nonetheless inconsistent. Finally, it is never possible to anticipate all fault modes of a system. Since alarm ranges are predefined during a design analysis in which fault modes of the system are enumerated, a new fault behavior may not trigger a predefined alarm at all.

The other traditional monitoring technique of discrepancy detection, normal behavior is obtained by simulating a model of the system being monitored. This approach avoids the limitations of the limit sensing approach. Normal behavior is fully context-sensitive, being derived via simulation from current state information. And in principle, new faults can be detected because there is no closed-world assumption based

on previously and statically enumerated known faults. However, the model-based approach of discrepancy detection has its own limitations. The approach is only as good as the system model, and some behaviors, e.g., non-linear behaviors, cannot be modeled with high fidelity with existing techniques. In addition, normal system behavior typically changes with time, and the model must continue to evolve. The modeling process must continue both because any long-lived system will degrade, and because a system may be used for different purposes throughout its lifetime. A good example of the latter is the Voyager 2 mission and spacecraft which were, respectively, replanned and reprogrammed to compensate for lower light levels for the extended mission to Uranus and Neptune [12].

Noting the limitations of the existing monitoring techniques, we have developed an approach to monitoring which is designed to make the anomaly detection process more robust, specifically to reduce the number of undetected anomalies (false negatives). Towards this end, we introduce multiple anomaly models, seven in all, each employing a different notion of “normal” behavior. Two of the seven are based on the traditional monitoring techniques. For each anomaly model there is a sensor importance measure. These measures determine why, at a particular moment, one sensor may be more worthy of operator attention than another. The anomaly models and sensor importance measures are based on concepts from model-based reasoning, statistics, and information theory. While any single one of these sensor importance measures may fail to detect an anomaly, we will demonstrate how collectively, they provide a finer safety net than the traditional monitoring techniques.

During each timestep all sensors are scored according to these sensor importance measures. The scores are used as a basis for an ordering on the sensors. See Figure 1,

These sensor scoring measures are partitioned into two broad categories. The first set – empirical methods – rely on current and historical data to determine importance, and take a purely local view of sensors. These measures include surprise, alarm, anticipate alarm, and value change. The alarm measure is based on the traditional notion of limit sensing. Measures from the second set – model-based methods – employ a model of the system and an event-driven simulation capability to reason about current and expected future system performance to determine sensor importance. These




Figure 2: Surprise of different sensor values and frequency distributions.

measures include *deviation*, *sensitivity*, and *cascading alarms*. The *deviation* measure is based on the traditional notion of discrepancy detection. The last two measures, *sensitivity* and *cascading alarms*, perform explicit causal reasoning on the model of a system.

This paper will concentrate on our n-odcl-based monitoring methods, providing most detail on our use of event-driven simulation and causal reasoning techniques. More detail on our empirical monitoring methods can be found in [8].

2.1 Empirical Anomaly Detection Methods

In this section, we briefly describe the empirical methods that we use to determine, from a local viewpoint, when a sensor is reporting anomalous behavior. There are four empirical sensor importance measures: *surprise*, *alarm*, *anticipate alarm*, and *value change*. These measures use knowledge about each individual sensor, without knowledge of any relations among sensors.

Surprise

An appealing way to assess whether current behavior is anomalous or not is via comparison to past behavior. This is the essence of the *surprise* measure. It is designed to highlight a sensor which behaves other than it has historically. Specifically, *surprise* uses the historical frequency distribution for the sensor in two ways: To determine the likelihood of the given current value of the sensor (*unusualness*), and to examine the relative likelihoods of different values of the sensor (*informativeness*). It is those sensors which display unlikely values when other values of the sensor are more likely which get a high *surprise* score. *Surprise* is not high if the only reason a sensor's value is unlikely is that there are many possible values for the sensor, all equally unlikely. See Figure 2.

The *informativeness* component of the *surprise* measure provides the key to detecting a subtle anomaly which is missed by both limit sensing and discrepancy detection, as will be discussed below.

Alarm

Alarm thresholds for sensors, indexed by operating mode, typically are established through an off-line analysis of system

design. The notion of *alarm* in SELMON extends the usual one bit of information (the sensor is in alarm or it is not), and also reports how much of the alarm range has been traversed. Thus a sensor which has gone deep into alarm gets a higher score than one which has just crossed over the alarm threshold.

Alarm Anticipation

The *alarm anticipation* measure in SELMON performs a simple form of trend analysis to decide whether or not a sensor is expected to be in alarm in the future. A straightforward curve fit is used to project when the sensor will next cross an alarm threshold, in either direction. A high score means the sensor will soon enter alarm or will remain there. A low score means the sensor will remain in the nominal range or emerge from alarm soon.

Value Change

A change in the value of a sensor may be indicative of an anomaly. In order to better assess such an event, the *value change* measure in SELMON compares a given value change to historical value changes seen on that sensor. The score reported is based on the proportion of previous value changes which were less than the given value change. It is maximum when the given value change is the greatest value change seen to date on that sensor. It is minimum when no value change has occurred in that sensor.

2.2 Model-Based Anomaly Detection Methods

Although many anomalies can be detected by applying anomaly models to the behavior reported at individual sensors, the most robust monitoring requires reasoning about interactions occurring in a system and detecting anomalies in behavior reported by several sensors.

The three model-based sensor importance scores in SELMON are *deviation*, *sensitivity*, and *cascading alarms*. While *deviation* only requires that some form of simulation be available for generating sensor value predictions, *sensitivity* and *cascading alarms* require the ability to simulate and reason with a causal model of the system being monitored.

In the next section, we describe the use of cvcml-driven simulation and explicit causal reasoning to support model-based monitoring methods. Towards this end, we include




Figure 3: A generic causal graph.

some detail on two tools we have built within the SELMON project: our EDSE event-driven simulator which accepts a causal model [14], and EIDSEL, our causal modeling language [21].

Deviation

The *deviation* measure is our extension of the traditional method of *discrepancy* detection. As in discrepancy detection, comparisons are made between predicted and actual sensor values, and differences are interpreted to be indications of anomalies. This raw discrepancy is entered into a normalization process identical to that used for the *value change* score, and it is this representation of relative discrepancy which is reported. The *deviation* score for a sensor is minimum if there is no discrepancy and maximum if the discrepancy between predicted and actual is the greatest seen to date on that sensor.

2.2.1 Causal Analysis

Many forms of reasoning relevant to monitoring can be supported by a *causal* model of a system. For the purposes of this discussion, we define a causal model to be one which provides explicit information about the quantities of a system, the structural connections through which quantities interact, and the mechanisms which govern how behavior propagates from one part of the system to another.

With this kind of information, it is possible to reason about: How faults may manifest in anomalies at sensors which are causally downstream and physically distant from the location of a fault. How the information reported at one sensor may be implicit in another, and hence redundant. How several sensors may be in an inconsistent state although each is within its nominal range and reporting a predicted value. How an apparently benign current state may contain the seeds of highly undesirable behavior about to manifest in the near future. All of these ideas on how to employ causal reasoning for monitoring purposes are being explored in the SELMON project. The most fully developed idea at this time is that of using a causal model to reason about future behaviors. In SELMON, this reasoning is embodied in the two causal measures *sensitivity* and *cascading alarms*.

Sensitivity measures the potential for a large global perturbation to develop from current state. *Cascading alarms* measures the potential for an alarm sequence to develop from current state. Both of these anomaly measures use a causal simulator to generate predictions about future states of the system, given current state. Current state is taken to be de-

finied by both the current values of system parameters (not all of which may be sensed) and the latent events already resident on the simulator agenda. The measures assign scores to individual sensors according to how the system parameter corresponding to a sensor participates in, or influences, the predicted global behavior. Roughly speaking, a sensor will have a high *sensitivity* score when behavior originating at that sensor causes a large amount of change elsewhere in the system. A sensor will have a high *cascading alarms* score when behavior originating at that sensor causes a large number of alarms elsewhere in the system.

Figure 3 shows a generic causal graph. System parameters, some of which will have associated sensors, are represented by nodes in this graph. Arcs represent the mechanisms which determine if and how system parameters affect other system parameters. There are two node subsets of interest relative to a given system parameter, or quantity node Q . These are $Causes_Q$, the set of quantity nodes upstream from Q in the directed causal graph, and $Effects_Q$, the set of quantity nodes downstream from Q in the directed causal graph. In computing a sensor's *sensitivity* or *cascading alarms* score, we are interested only in behavior which "passes through" the quantity corresponding to the given sensor. The only events processed by the simulator are those involving the quantity itself, its causes, or its effects. Initially, the simulator only processes events due to the quantity itself or its causes. However, once an event at the quantity itself is processed, the simulator also processes events among the effects of the quantity. These events are said to derive from the quantity. An event at a quantity which is not reachable from Q is not processed by the simulator, even if it is on the simulator agenda, for such behavior cannot have been influenced by the quantity of interest. The only simulated events which contribute to a sensor's *sensitivity* or *cascading alarms* score are those involving the quantity itself or its (sensed) effects. Simulation proceeds until a stated future time is reached. Currently, we assume there is no feedback, i.e., the causal graph is a directed acyclic graph.

Sensitivity

The algorithm used to compute a sensor's *sensitivity* score is as follows:

Given current system state

(quantity values at t_0 and queued events)

For each quantity Q :

Retrieve *Causes* $_Q$

Retrieve *Effects* $_Q$

Simulate only those events originating or deriving from a quantity $g \in \{Q\} \cup \text{Causes}_Q$

Collect all events at a sensed quantity $Q_e \in \text{Effects}_Q$ occurring at time $t \leq t_0 + \Delta t$:

Let

$$\text{Sensitivity}(Q) = \sum \frac{|Q_e(t) - Q_e(t_0)|}{\text{Max_Change}(Q_e)}$$

t_0 is the current time. Δt is the maximum time forward 10 simulate. $\text{Max_Change}(Q_e)$ refers to the largest historical value change seen on a given sensor. The values which are summed are the ratios between the maximum predicted value change for a given sensor Q_e and the maximum historical value change for that sensor. The *sensitivity* score is maximum when every sensor causally downstream from a given sensor is predicted to exhibit its largest value change seen to date. Thus the maximum score for a given sensor is equal to the number of sensors causally downstream from that sensor in the causal graph. The *sensitivity* score is minimum when no changes are predicted to occur causally downstream from a given sensor. All *sensitivity* scores are zero if the system is in a perfectly steady state.

Note that for a given set of predicted, causally related events, a sensor's score is monotonically higher the closer it is to the causal source of that global behavior. This information may be useful, for example, in aiding operators in selecting a single control action with the greatest impact.

Cascading Alarms

The algorithm used to compute a sensor's *cascading alarms* score is as follows:

Given current system state

(quantity values at t_0 and queued events)

For each quantity Q :

Retrieve *Causes* $_Q$

Retrieve *Effects* $_Q$

Simulate only those events originating or deriving from a quantity $g \in \{Q\} \cup \text{Causes}_Q$

Collect all events at a sensed quantity $Q_e \in \text{Effects}_Q$ occurring at time $t \leq t_0 + \Delta t$:

Let

$$\text{CascadingAlarms}(Q) = \sum \text{In_Alarm}(Q_e(t))$$

$\text{In_Alarm}(Q_e(t))$ is 1 if the quantity Q_e is in an alarm range at simulated time t ; otherwise it is 0. The *cascading alarms* score is maximum when every sensor causally downstream from a given sensor is predicted to be in alarm. Thus the maximum score for a given sensor is equal to the number of sensors causally downstream from that sensor in the

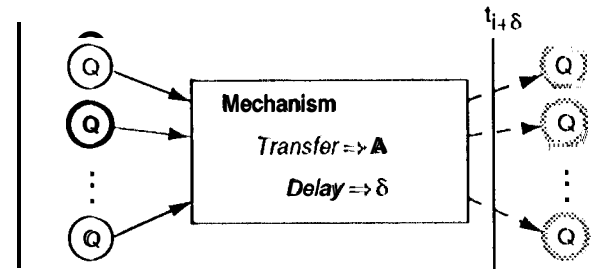


Figure 4: Causal simulation in EDSE: A mechanism evaluated at time t_i will propagate its value at $t_i + \delta$.

causal graph. The *cascading alarms* score is minimum when no alarms are predicted to occur causally downstream from a given sensor.

Note that the *cascading alarms* score is designed to give monotonically higher scores to sensors that are closer to the causal source of an alarm sequence. This kind of information may be extremely valuable to an operator attempting to interpret a situation where multiple alarms appeared between one polling of the sensors and the next.

2.3 Causal Simulation and Modeling

In this section, we briefly describe the EDSE event-driven causal simulator and the EDSEL causal modeling language which support the SELMON system.

Causal models are an abstraction for describing the behavior of a system by representing the physical processes occurring within the system as discrete functions to be evaluated.

A *causal model* characterizes a physical system in terms of state variables and causal influence relations among the variables. In the *causal graph* defined by the state variables and influence relations, changes in any variable may be propagated to other variables through the influence relations. Causal simulation is the process of tracking changes in variables and propagating them to other variables through influence relations, thereby producing a new set of changes. Implicit in the notion of causality are the concepts of *event* and *causal time*: Events comprise changes in state variables due to a specific influence relation and with respect to a specific moment in time. Thereby, causal time moves forward due to delays in the propagation of changes in the causal model.

In our causal modeling language, called EDSEL, causal models are represented in terms of these primitives. Causal time is expressed as a monotonically increasing sequence of integer-valued *instants*. Each state variable is denoted by a *quantity*. Each influence relation is described by a *mechanism* which encapsulates a set of input quantities, a set of output quantities, along with *transfer* and *delay* functions. During simulation in EDSE, the transfer function is evaluated with the mechanism input quantities to produce a value that will be propagated to each of the mechanism output quantities. The delay function is simultaneously evaluated to produce an offset from the current simulation time. Value propagation takes place, at the relative time returned by the delay function. Figure 4 visualizes mechanism evaluation. Simulation continues until a user-specified time limit is reached or the model reaches quiescence.

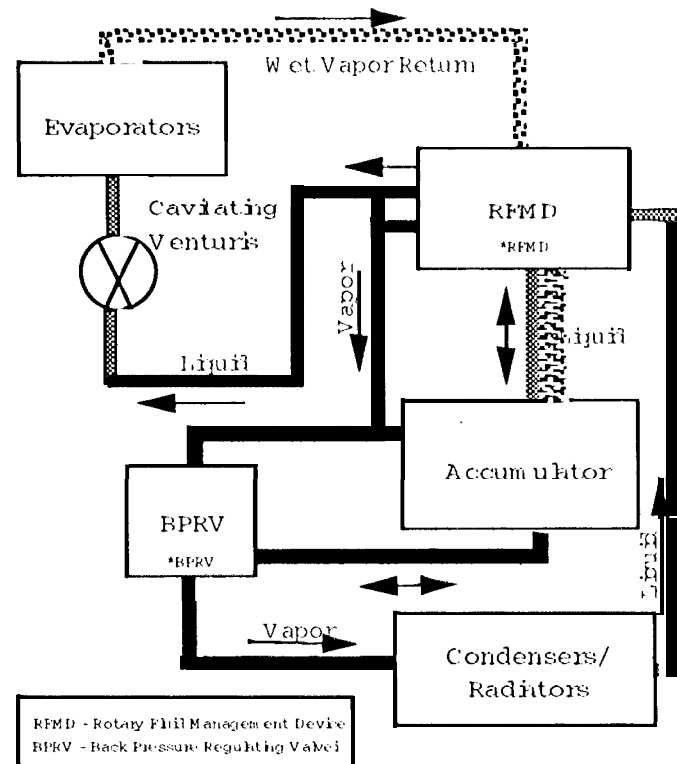


Figure 5: The External Active Thermal Control System (EATCS) of Space Station Freedom.

We adopt a component-centered approach to causal modeling. This modeling approach distinguishes the constituents of the system being modeled as *components* and the structural *connections* among them. A physical component is intuitively defined as a discrete element of a physical system. By analogy, a causal component is an entity that encapsulates the behavior of the corresponding physical component as a whole.

A connection is an abstraction which characterizes the interaction pathways between components. A connection within a physical system enables material or information flow between components. Similarly, causal connection allows quantity values to be propagated between components.

2.3.1 Modeling the SSF EATCS

We have commenced a modeling effort to build a causal model of the SSF EATCS, using our G2-based tool called MESA [31] which generates the EDSEL representation of a causal model from a higher-level abstraction of a component-centered model together with the mechanisms which represent the physical behavior of the components. The MESA-generated model is then loaded into EDSEL to perform an event-driven simulation of the EATCS.

Figure 5 represents a simplified schematic of the EATCS, showing all of the major components. The Rotary Fluid Management Device (RFMD) pumps liquid ammonia to the Evaporator loop. Flow control from the RFMD to each of the evaporators is provided by cavitating venturis. In a normal venturi with fixed inlet conditions, reducing the exit pressure results in an increased flowrate. The TCS cavitating venturi passively provide flow stability even with downstream pressure variations resulting from heat load changes. A two-phase

fluid mixture of liquid and vapor is generated at each evaporator by the particular heat load being serviced. The two-phase mixture is returned to the RFMD, which separates the liquid from the vapor by centrifugal force. The RFMD pumps the vapor to the radiators, which condense the fluid to a subcooled state by rejecting heat to space. The Back Pressure Regulating Valve (BPRV) maintains setpoint temperature by controlling system pressure. RFMD fluid level changes are due primarily to varying heat loads and are handled by an accumulator with an internal metal bellows separating the liquid and vapor sides.

The component model schematic of a simplified EATCS evaporator loop is shown in Figure 6. Two cavitating venturis, CV1 and CV2, regulate flow to two evaporators, EV1 and EV2, operating in parallel. The RFMD is modeled as a flow and pressure source for the evaporator loop, and divide and jointees DV1 and JT1 represent the flow junctions in the loop. H1 and H2 represent the heat sources for the evaporators. The modeler creates instances of these components from a component library which contains the definitions of component classes, their input and output quantities, and the mechanisms which encapsulate the transfer and delay functions with the quantities. An example of a component definition for the cavitating venturi is shown in Figure 7. This representation is built by a developer in MESA for the component library, and the instantiation of the actual underlying component quantities, transfer and delay functions, and the inter-component connections is handled automatically by MESA when the component model schematic is built. Thus, the modeler is freed from details of building the causal model directly in EDSEL.

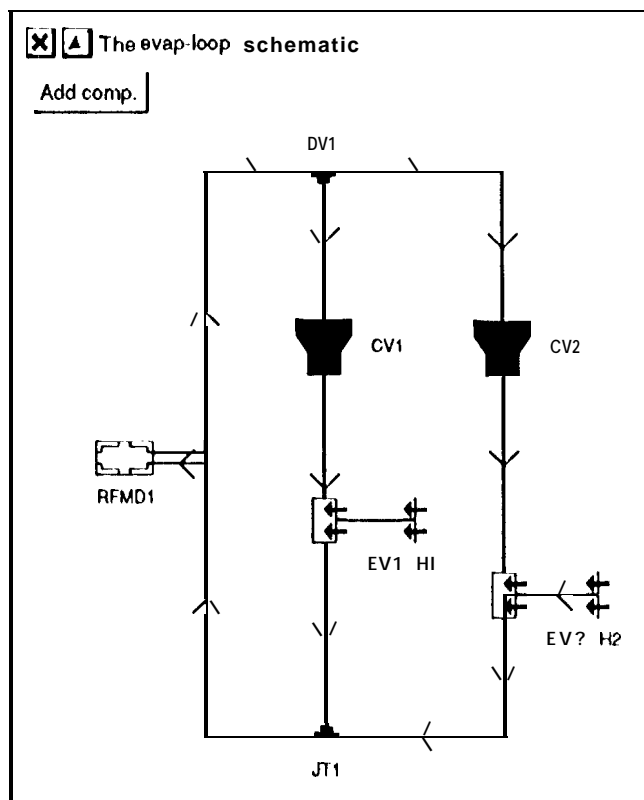


Figure 6: The Evaporator Loop of SSF EATCS.

The component definition of the venturi as shown in Figure 7 shows the flow of input quantities through transfer functions to output quantities for each state variable that describes a cavitating venturi. These transfer functions are based on the laws of physics, fluid dynamics and thermodynamics, and contain a complete mathematical description of the steady state behavior of the device. Each transfer function by definition may have multiple inputs, but produces a single output. In cases where intermediate calculations must be made to determine an output quantity, such as for saturation pressure and liquid density at the venturi inlet temperature used for flow rate determination, mechanisms are chained together to produce the final result.

Using these generic modeling techniques, a causal model of the simplified EATCS evaporator loop has been generated in EDSEL, and simulations were executed successfully in EDSE. We are currently working to model the EATCS radiator loop and the remainder of the RFMD and BPRV. The next step will be to incorporate sensor objects into the model to support the model-based sensor importance measures in SELMON.

3 Performance Evaluation

The SELMON sensor importance measures are dynamically computed each time the sensors are polled. In order to assess whether SELMON usefully focuses operator attention while performing robust anomaly detection, we performed the following experiment: We evaluated SELMON performance in selecting critical sensor subsets specified by an SSF ECLSS domain expert, sensors seen by that expert as useful in under-

standing episodes of anomalous behavior in actual historical data from ECLSS testbed operations.

The experiment asked the following two specific questions: How often did SELMON place a "critical" sensor in the top half of an overall ordering on the sensors? and How does SELMON performance compare to traditional monitoring practice?

We used two alternate ways of combining the scores from the seven sensor importance measures to arrive at an overall ordering: 1) the individual scores were composed arithmetically (see [81] for details) and 2) the maximum individual score was taken. We also ran SELMON with only the *alarm* measure active, to represent traditional monitoring.

The performance of a random sensor selection algorithm would be expected to be about 50.070: any particular sensor would appear in the top half of the sensor ordering about half the time. Table 1 shows the results of our experiment. The first column identifies one of the anomaly episodes specified by the domain expert. The second column shows the overall "hit" rate for each episode using the *alarm* measure only, i.e., the percentage of time SELMON placed the given sensor in the top half of the sensor ordering generated by the *alarm* measure. The third column shows the hit rate using arithmetic composition of the scores. The fourth column shows the hit rate when the maximum score is used. Finally, the fifth column shows the hit rate when either way of combining the scores is used.

EPISODE	Alarm	Composed	Max	Either
High Flow Rate	97.8	94.5	83.5	94.5
Sensor Malfunction	100.0	100.0	100.0	100.0
Unibed Loading	48.9	56.6	90.6	90.6
Pre-Heater Off	100.0	100.0	81.1	100.0
Emergency Shutdown	100.0	100.0	100.0	100.0
Pressure Fluctuations	98.6	100.0	90.1	98.6
High Pressure	92.5	95.5	91.0	92.5
All	76.3	79.8	88.1	95.1

Table 1: SELMON Performance at selecting critical sensor data.

These results show that SELMON performs much better than random at replicating the attention focusing of an ECLSS domain expert. More to the point, when the maximum sensor importance measure is used for each sensor, SELMON performs considerably better than the traditional practice of limit sensing. Finally, when both methods of combining the individual scores are available (noting that, at the current time, we do not have an automated technique for determining in context which method is more appropriate, and such determination would be left to the operator), SELMON performance is quite notable, detecting actual anomalies 95% of the time.

These results lend credibility to our premise that the most effective monitoring system is one which incorporates several models of anomalous behavior, and takes multiple views of sensor importance. The best, most experienced mission operators are already remarkably effective at knowing when something is going wrong on a space platform. Our aim is to offer a more complete, more robust set of techniques for anomaly detection, to make mission operators even more effective, or to provide the basis for an automated monitoring capability.

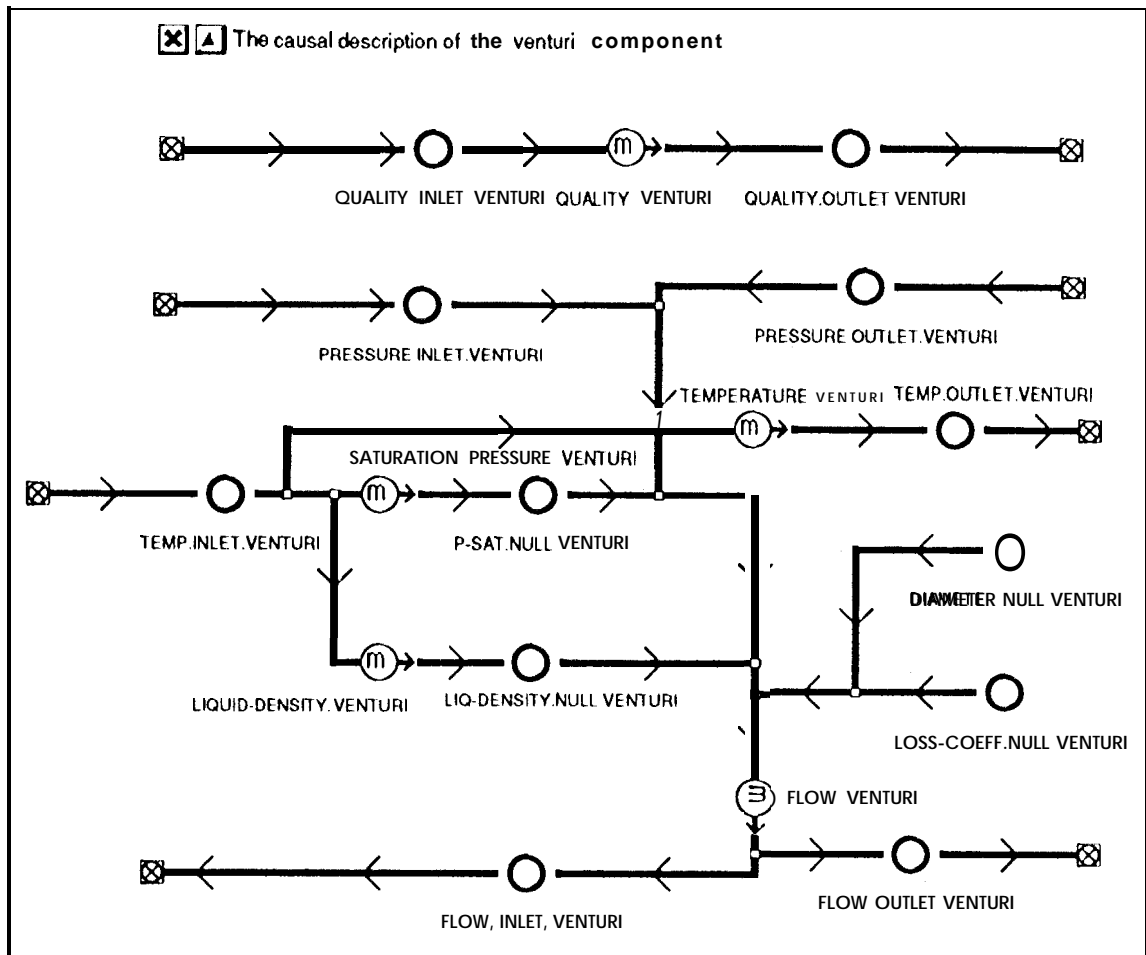


Figure 7: The Venturi component of SSF'EATCS.




Figure 8: A failure on the water side of SSF ECLSS.

We plan to extend this empirical analysis of SELMON performance in two directions: 1) using both the alarm and deviation measures as the baseline for performance comparison, to better represent traditional monitoring techniques, and 2) to examine the performance of the causal measures sensitivity and cascading alarms separately, possibly in artificial domains, where the topology of the causal graph and the mechanisms can be specified to better analyze how these measures perform under different situations of undesirable behavior propagation.

Besides examining SELMON performance at the statistical level, we have also examined individual examples of successful anomaly detection in SELMON. The following example illustrates how SELMON highlights a subtle manifestation of an anomaly which the traditional monitoring techniques fail to detect.

During an episode when the ECLSS testbed pre-heater failed (see the schematic in Figure 8), system pressure, which normally oscillates within a known range, became more stable (see Figure 9). This "abnormally normal" behavior is not detected by limit sensing because the system pressure remains firmly in the nominal range; nor by discrepancy detection, for the predicted value is roughly the average of the values over time, this oscillating behavior being unmodelled. However, the SELMON informativeness component of the surprise measures rises during this episode. Informativeness rises when the frequency distribution across the range of sensor values moves away from a flat distribution towards a "spike" distribution (see Figure 2). A suddenly stable system pressure results in one of the value ranges for system pressure beginning to dominate the frequency distribution (see Figure 10). SELMON provides the means of detecting and reasoning about this kind of subtle anomaly.




Figure 9: ECUS System Pressure During Abnormally Normal Episode.




Figure 10: Informativeness of ECLSS System Pressure During Abnormally Normal Episode.

4 Towards Applications

We are working to apply the SELMON techniques in an advanced monitoring and diagnosis prototype for the SSF External Active Thermal Control System (EATCS). The Thermal Control System Automation Project (TCSAP) [11] is developing a Jmow/cdgc-based system to perform Fault Detection, isolation and Recovery (FDIR) on the EATCS. The TCSAP software is a hybrid advanced automation system implemented in a commercial rzczl-time expert system shell called G2. It uses a combination of conventional programming, rule-based technology, and model-based reasoning to provide extensive diagnostic power and flexibility beyond the capabilities based on alarm thresholds and limit sensing. Within TCSAP, the SELMON empirical sensor importance measures described earlier have been implemented, and testing and evaluation have begun in ground-based Thermal Test Bed runs at NASA Johnson Space Center (JSC).

Also at NASA JSC, we are exploring applications of the SELMON approach to Space Shuttle (STS) operations. In particular, we are examining the Environmental Emergency and Consumables Management (EECOM) subsystem of STS. This is the life support system for STS, the analog of ECLSS for SSF.

The seven anomaly measures in SELMON are designed to be applied separately. This feature allows a SELMON-based application to be prototyped and evaluated in an efficient manner. As long as sensor data is available, on-line or historical, the four empirical (non-model-based) anomaly measures can be immediately tested. If any form of simulation is available, the deviation measure also can be tested. It is only the two causal measures which must await for the causal modeling process to complete. Thus SELMON avoids a common limitation of AI-based systems development: that no delivery is possible, even a prototype system, before a large knowledge engineering or modeling investment is made.

5 Future Work

An unresolved area in SELMON is developing a well-founded method for utilizing all the sensor importance measures. We have working concepts for how to compose the individual measures into a total sensor importance score, but we consider this an area for further work. A theoretical or empirical analysis may provide insight on the most appropriate composition technique, one that is tailorable to different applications.

We recognize that an important component of the SELMON approach is the ability to provide explanations or interpretations of why a particular sensor has been highlighted and is more worthy of operator attention than other sensors. Other future work in the SELMON project will complement existing sensor ordering and anomaly detection capabilities with model-based capabilities for characterizing anomalies by their temporal and spatial extent, and focusing attention according to multiple viewpoints (causal priority, proximity to control points, potential for irreversible damage, etc.).

In related work, we are also investigating the problem of sensor placement during design [4].

6 Summary

We are developing techniques to support real-time monitoring through sensor selection, the moment to moment focusing

of attention on a subset of the available sensor data. Sensor selection is based on a set of importance criteria using different models of what constitutes an anomaly. The computational realizations of these importance criteria draw on concepts from model-based reasoning, statistics, and information theory. Experimental results show that our sensor selection techniques are effective at highlighting the sensors deemed critical by a domain expert for understanding actual anomalous episodes from the Space Station Freedom Environmental Control and Life Support System testbed. These results also suggest that a monitoring system which employs multiple models of anomalous behavior is more effective than one based on the traditional monitoring concepts of alarm thresholds and discrepancies.

7 Acknowledgements

Others who have worked recently on the SELMON project include Steve Chien, Daniel Clancy, Usama Fayyad and Harry Porta.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] K. J. Abbott, "Robust Operative Diagnosis as Problem Solving in a Hypothesis Space," *7th National Conference on Artificial Intelligence*, St. Paul, Minnesota, August 1988.
- [2] L. Charest, Jr. and R. Faltischo, "Building Causal Models for SELMON: A User's Manual for the EDSEL Modeling Language," JPL Publication, in preparation.
- [3] L. Charest, Jr., N. Rouquette, R. Doyle, C. Robertson, and J. Wyatt, "MESA: An Interactive Modeling and Simulation Environment for Intelligent Systems Automation," submitted to *Hawaii International Conference on System Sciences*, Maui, Hawaii, January 1994.
- [4] S. A. Chien, R. J. Doyle, and N. F. Rouquette, "A Model-based Reasoning Approach to Sensor Placement for Diagnosability," *2nd International Workshop on the Principles of Diagnosis*, Milan, Italy, October 1991.
- [5] D. DeCoste, "Dynamic Across-Time Measurement Interpretation," *8th National Conference on Artificial Intelligence*, Boston, Massachusetts, August 1990.
- [6] R. J. Doyle, J. J. Atkinson, and R. S. Doshi, "Generating Perception Requests and Expectations to Verify Joint Execution of Plans," *5th National Conference on Artificial Intelligence*, Philadelphia, Pennsylvania, August 1986.
- [7] R. J. Doyle, D. Berleant, L. Falcone, and U. Fayyad, "Selective Simulation and Selective Sensor Interpretation in Monitoring," *Computers in Aerospace*, American Institute of Aeronautics and Astronautics, October 1989.
- [8] R. J. Doyle, S. A. Chien, U. M. Fayyad, and E. J. Wyatt, "Focused Real-time Systems Monitoring Based on Multiple Anomaly Models," *7th International Workshop*

on *Qualitative Reasoning*, Eastsound, Washington, May 1993.

- [9] D. L. Dvorak and B. J. Kuipers, "Model-based Monitoring of Dynamic Systems," *11th International Conference on Artificial Intelligence*, Detroit, Michigan, August 1989.
- [10] D. L. Dvorak, "Monitoring and Diagnosis of Continuous Dynamic Systems using Semiquantitative Simulation," Ph.D. dissertation, The University of Texas at Austin, May 1992.
- [11] T. Hill, W. Morris, and C. Robertson, "Implementing a Real-time Reasoning System for Robust Diagnosis," *6th Annual Workshop on Space Operations Applications and Research*, August 1992.
- [12] R. P. Laeser, W. I. McLaughlin, and D. M. Wolff, "Engineering Voyager 2's Encounter with Uranus," *Scientific American*, 255(5):36-45, November 1986.
- [13] J. Muratore, T. Heindel, T. Murphy, A. Rasmussen, and R. McFarland, "Space Shuttle Telemetry Monitoring by Expert Systems in Mission Control," in *Innovative Applications of Artificial Intelligence*, J. Schorr and A. Rappaport (eds.), AAAI Press, 1989.
- [14] N. F. Rouquette, S. A. Chien, and L. K. Charest, Jr., "Event-driven Simulation in SELMON: An Overview of EDSE," JPL Publication 92-23, August 1992.
- [15] E. A. Scarf, J. R. Jamieson, and E. New, "Deriving Fault Location and Control from a Functional Model," *3rd IEEE Symposium on Intelligent Control*, Arlington, Virginia, 1988.